From Slow to Scalable: Strategies for Large-scale Mixed Integer Linear Programming Optimisation

Leena Aizdi - PhD Candidate
Department of Computer Science, University of Pisa
March 28, 2025

Agenda

01

Brief Introduction to MILP

02

Problem Definition

03

Current Research & Techniques

04

Conclusions & Future Prospects

What is Mixed Integer Programming

A Mixed Integer Linear Programming (MILP) formulation implies the presence of both continuous ($x \in \mathbb{R}$) and discrete ($y \in \mathbb{Z}$) variables.

$$\max \ \mathbf{c}'\mathbf{x} + \mathbf{d}'\,\mathbf{y} \quad \longleftarrow \text{Objective Function}$$

$$\left.\begin{aligned} s.t.\ \ \mathbf{A}\mathbf{x} + \mathbf{E}\mathbf{y} &= \mathbf{b} \\ \mathbf{x} &\geq 0 \\ \mathbf{y} &\in \mathbb{Z} \end{aligned}\right\} \text{Constraints}$$

(Bigi et al., 2024)

What is Mixed Integer Programming

For example:

A farmer want to **maximise profit** by allocating land for two crops:

- Wheat (can be planted in any fraction of an hectare, so it's a **continuous** variable).
- Corn (must be planted in whole hectares, so it's a **discrete** integer variable).

The **total land** available is **10 hectares**. The **profit per hectare** for wheat is **200$** and that for corn is **300$**.

MILP Model:

Decision Variables:

- x_W = Hectares of **wheat** planted (**continuous variable**).
- x_C = Hectares of **corn** planted (**integer variable**).

Objective Function (Maximize Profit):

$$\text{Maximize } Z = 200x_W + 300x_C$$

Constraints:

1. **Total Land Constraint:**

$$x_W + x_C \leq 10$$

(Total land used cannot exceed 10 hectares)

2. **Non-Negativity & Type Constraints:**

$$x_W \geq 0, \quad x_C \geq 0, \quad x_C \text{ is an integer}$$

What is Mixed Integer Programming

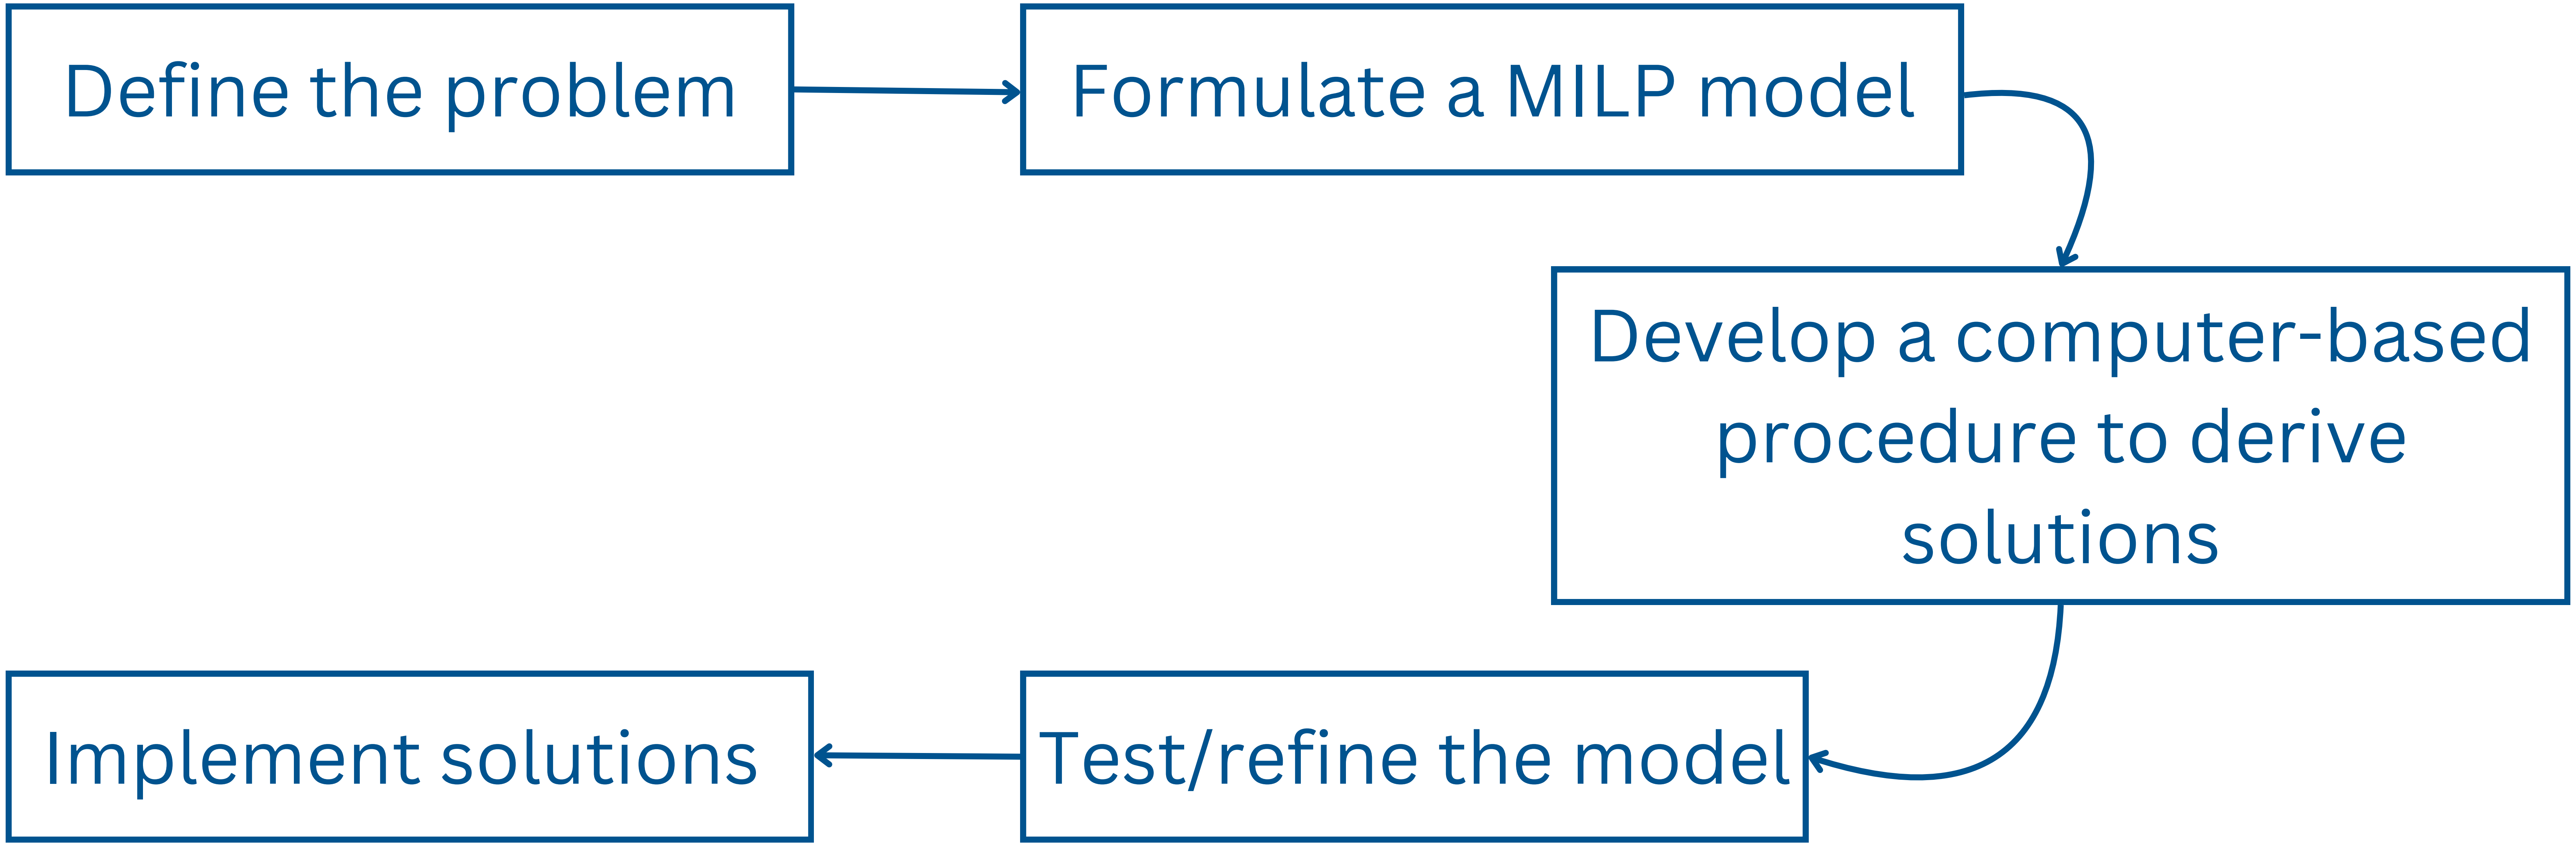

(Hillier & Lieberman, 2015)

What is the Problem?

As the size of the model grows, it becomes very difficult to find optimal or near-optimal solutions in a reasonable amount of time using over-the-shelf solvers.

Current Research & Techniques

Solutions include:

1. Exploiting Exact Decomposition
2. Using Heuristic Approaches
3. AI/ML Integrated Optimisation

Current Research & Techniques

Exploiting Exact Decomposition

Current Research & Techniques

Exploiting Exact Decomposition

(Müllerklein & Fontaine, 2025)

- The authors aim to design a decision-support model that helps optimise supply chain resilience in the face of transportation disruptions.
- Exploit **Benders Decomposition**

Current Research & Techniques

Exploiting Exact Decomposition

Benders Decomposition - The idea

Decompose the problem into:

- Master Problem: Focuses on complicated variables; A solution here provides initial values for these variables.
- Subproblem: Takes the fixed complicated variables from the master problem, determines how they interact with the continuous variables and sends feedback to the master problem in the form of "cuts" (new constraints).
    - Feasibility/Optimality Cuts

Current Research & Techniques

Exploiting Exact Decomposition

The problem is:

$$\text{Minimize} \quad \mathbf{c}^\top \mathbf{x} + \mathbf{d}^\top \mathbf{y}$$

Subject to:

$$A\mathbf{x} + B\mathbf{y} \geq \mathbf{b}$$

$$\mathbf{y} \geq 0$$

Where:

- $\mathbf{x} \in \mathbb{Z}^n$ (discrete)
- $\mathbf{y} \in \mathbb{R}^m$ (continuous)

1. **Master Problem (MP)**:

   The master problem (MP) involves the discrete variables $\mathbf{x}$. It is formulated as:

   $$\text{Minimize} \quad \mathbf{c}^\top \mathbf{x} + \theta$$

   Subject to:

   - $\mathbf{x} \in X$ (discrete set)
   - Benders **cuts** (which will be generated iteratively from the subproblem)

2. **Subproblem**:

   For a fixed $\mathbf{x}^b$ (from the master problem), solve the **subproblem** in $\mathbf{y}$:

   $$\text{Minimize} \quad \mathbf{d}^\top \mathbf{y}$$

   Subject to:

   $$B\mathbf{y} \geq \mathbf{b} - A\mathbf{x}^b$$

   $$\mathbf{y} \geq 0$$

   The subproblem is typically a **Linear Program (LP)** and provides the value $\Phi(\mathbf{x}^b)$ for each $\mathbf{x}^b$.

Current Research & Techniques

Exploiting Exact Decomposition

Further enhancements in Benders Decomposition (Müllerklein & Fontaine, 2025):

- Lower-bound lifting and valid inequalities: Strengthening the solution space to improve by incorporating problem-specific constraints to eliminate infeasible solutions early.
- Warm-start heuristics: Using initial feasible solutions to reduce computation time.
- Branch-and-Benders-cut: Combining branch-and-bound and Benders decomposition for better performance.

Results show that large instances can be solved to near-optimality, whereas a commercial solver does not find feasible solutions.

Current Research & Techniques

Using Heuristic Approaches

Current Research & Techniques

Using Heuristic Approaches

- Heuristics provide near-optimal solutions when finding the exact solution is too time-consuming or computationally expensive.
- They rely on simplified rules or logic rather than exhaustive calculations.
- Usually tailored to particular problems
- For example, greedy algorithms

Current Research & Techniques

Using Heuristic Approaches

An example Matheuristic:

Algorithm TD: Time horizon decomposition matheuristic

Divide the time horizon into Λ subperiods
for $\lambda = 1, \ldots, \Lambda$ **do**
    Solve subproblem λ
end for
Concatenate the subproblem solutions from 1 to Λ

(Lanza et al., 2023)

Current Research & Techniques

AI/ML Integrated Optimisation

Current Research & Techniques

Why do we need AI/ML Integrated Optimisation?

- Incorporating AI and ML can automate the traditional methods based on handcrafted heuristics, making the process more systematic.
- Many decisions in MILP solvers, such as branching, and node selection, are based on heuristics. AI/ML can help optimise these decisions by learning from data, potentially improving solver efficiency.
- Deep learning excels in high-dimensional spaces and can be leveraged to address the complexities of MILP problems effectively.

(Clautiaux & Ljubić, 2024)

Current Research & Techniques

AI/ML Integrated Optimisation

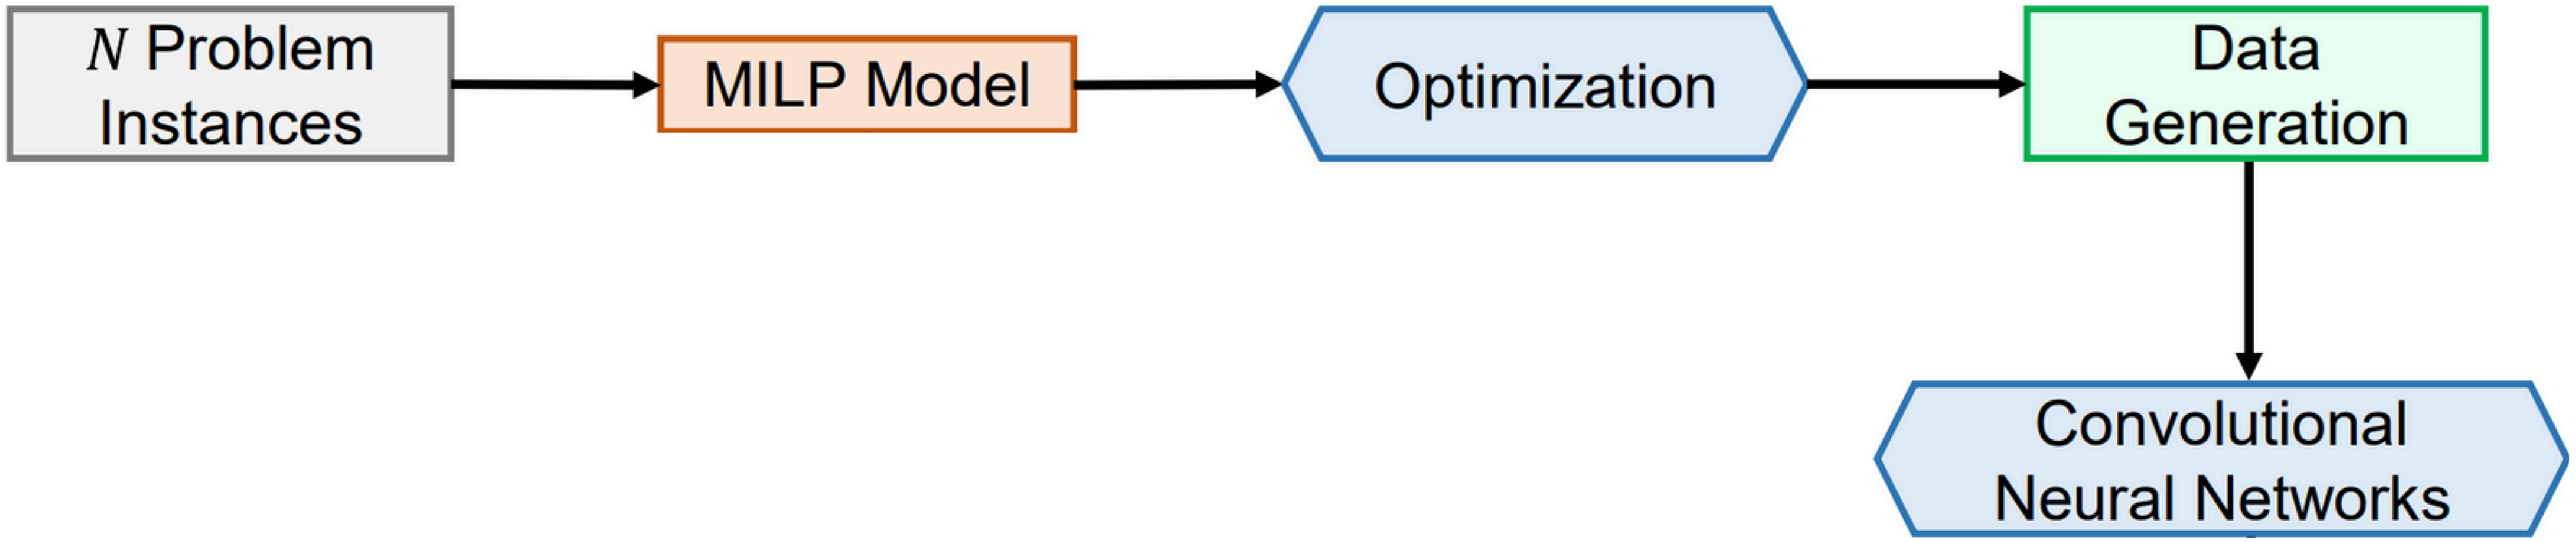

(Triantafyllou & Papathanasiou, 2024)

Current Research & Techniques

AI/ML Integrated Optimisation

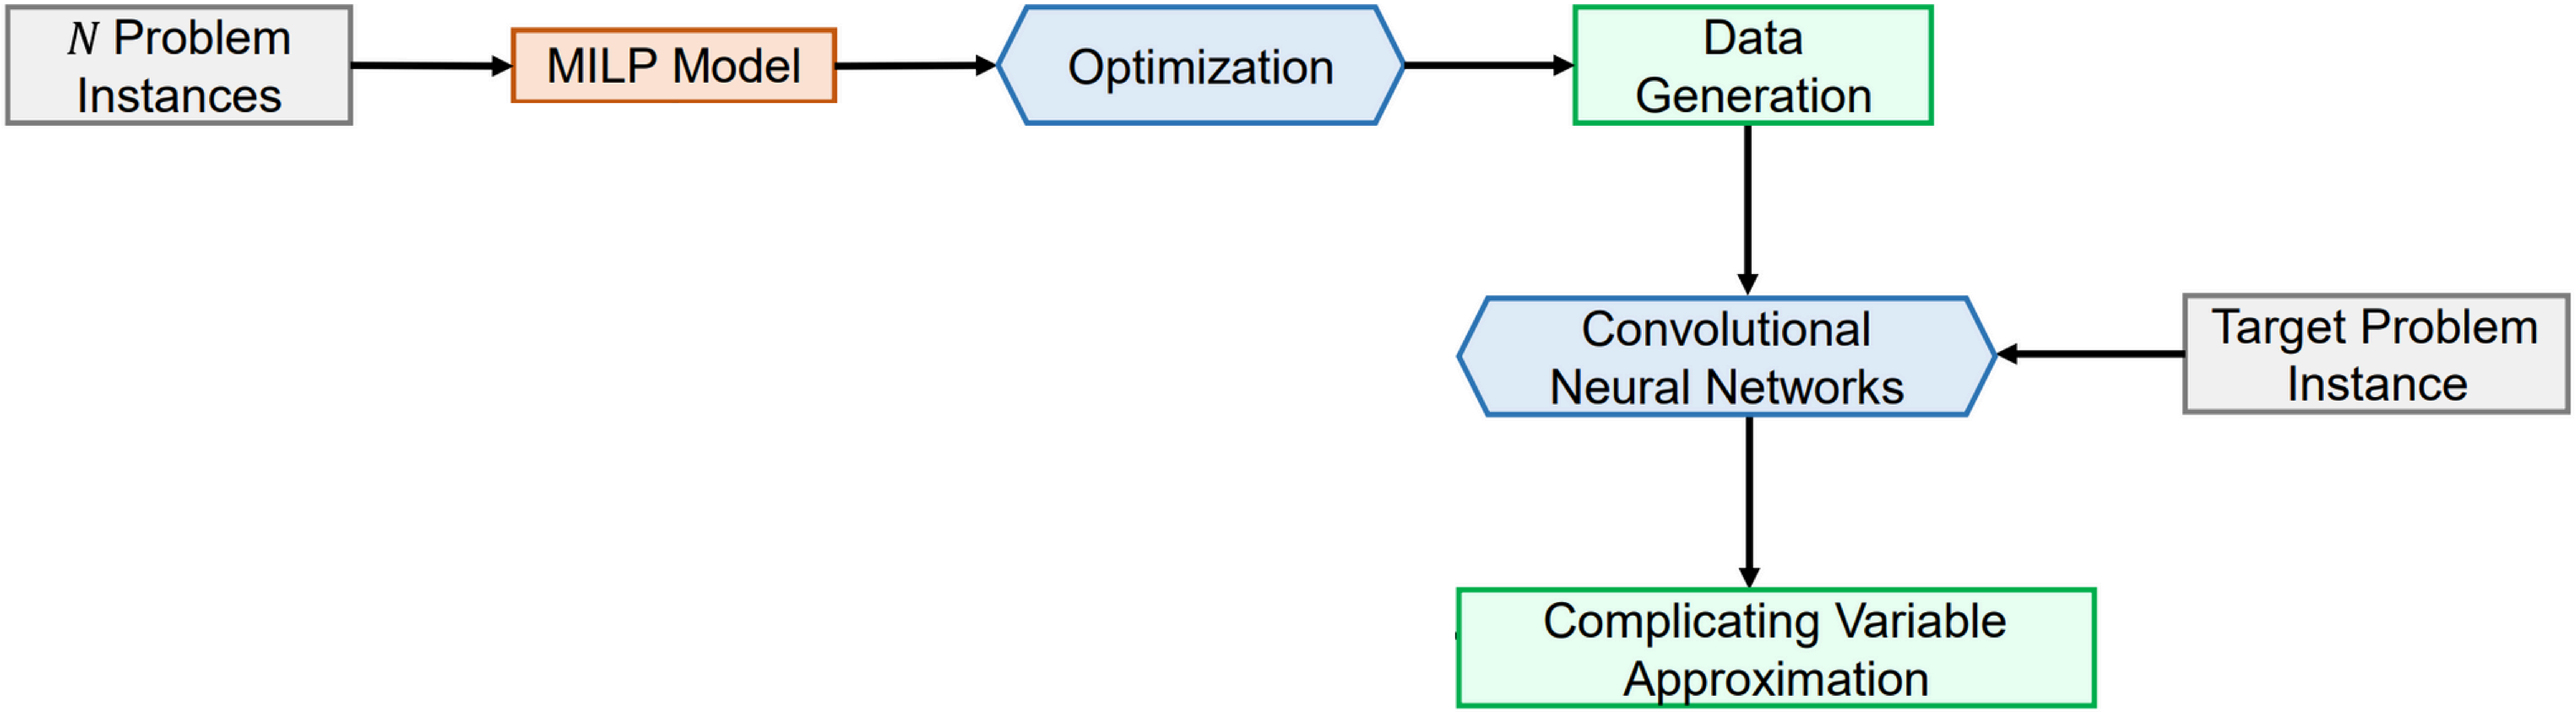

(Triantafyllou & Papathanasiou, 2024)

Current Research & Techniques

AI/ML Integrated Optimisation

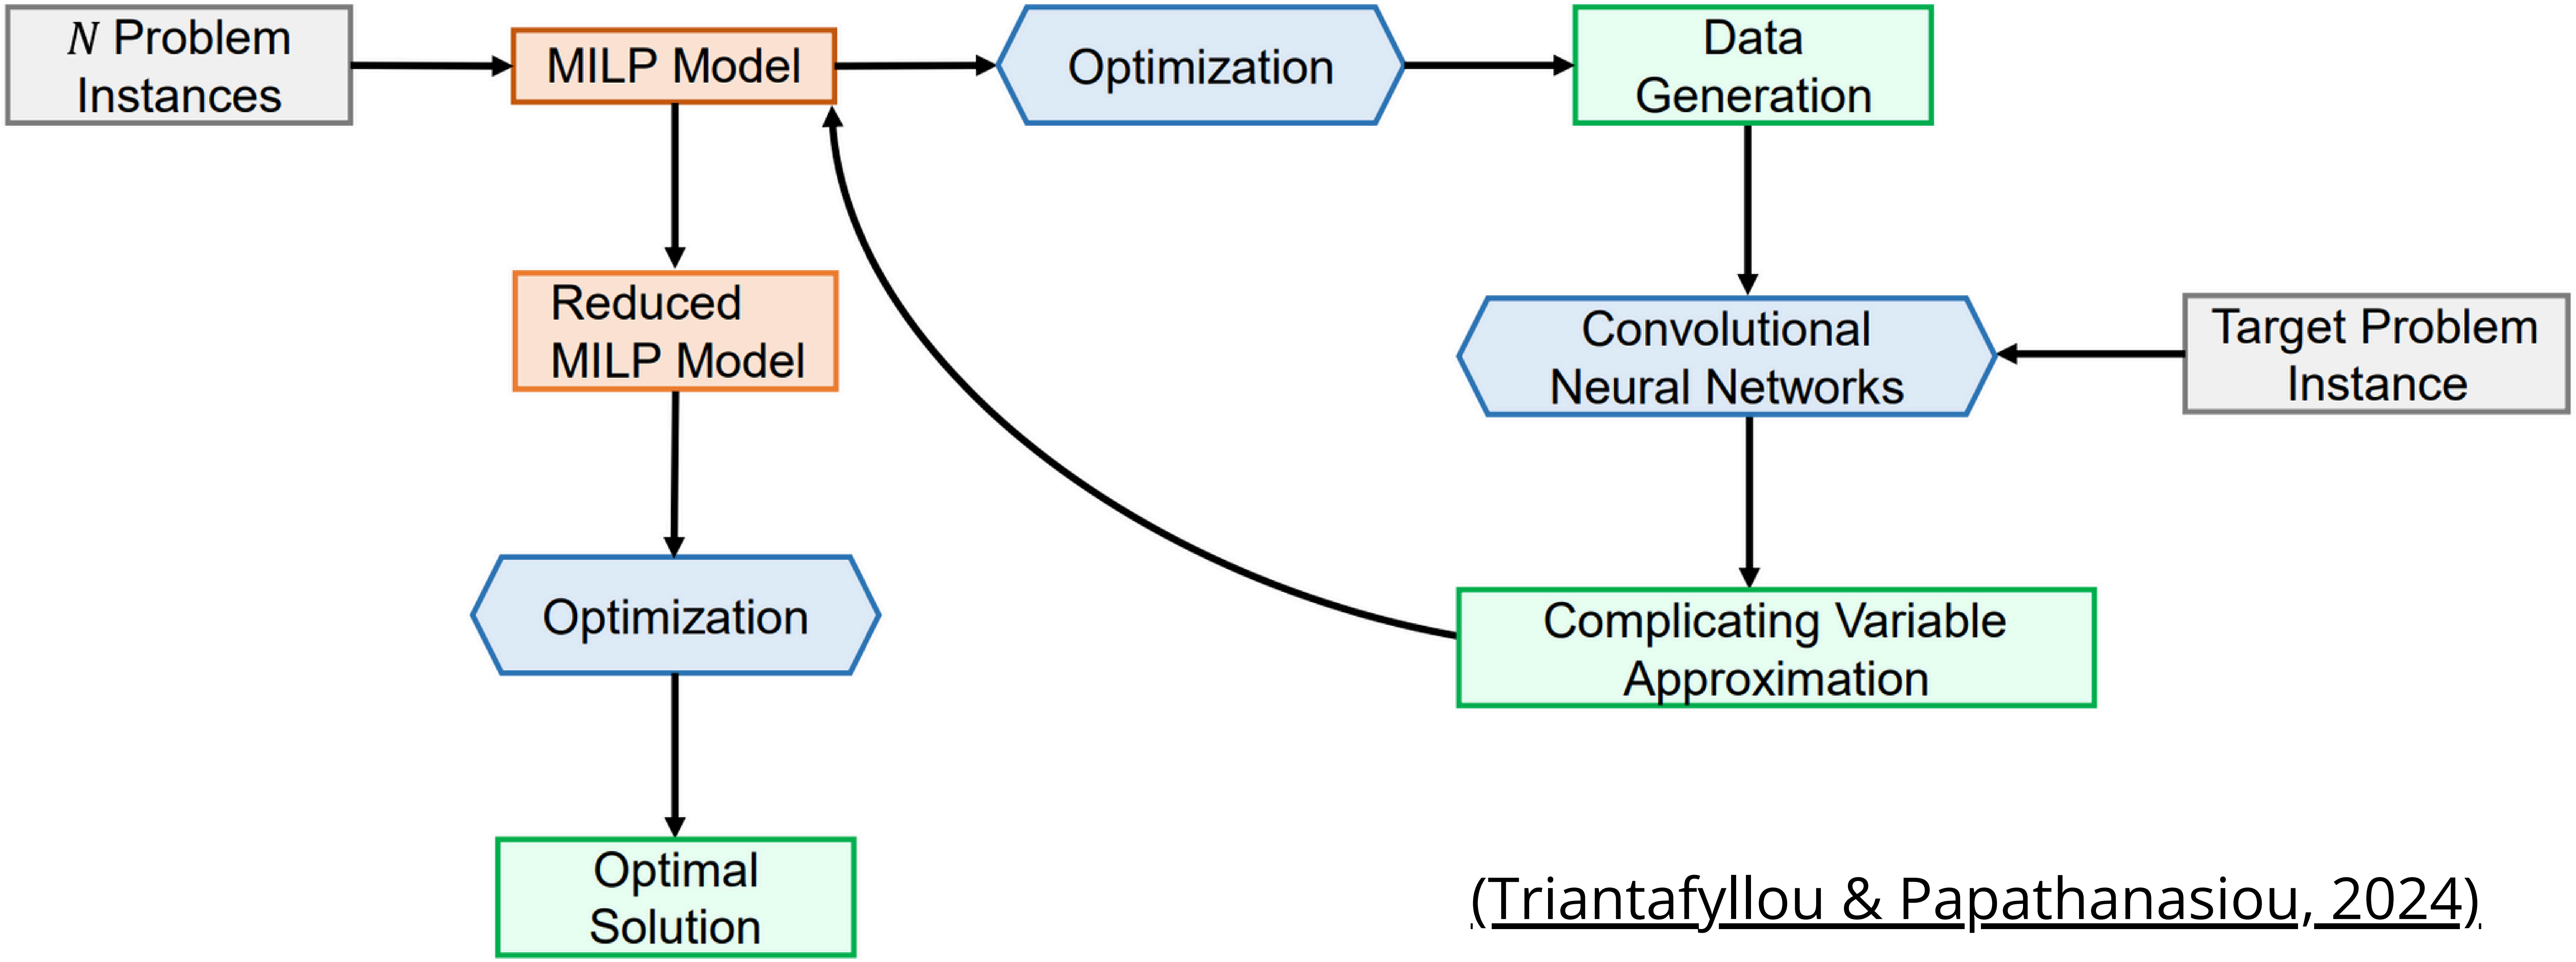

(Triantafyllou & Papathanasiou, 2024)

Current Research & Techniques

AI/ML Integrated Optimisation

Convolutional Neural Network (CNN) Training & Architecture:

- Dataset: Contains over 8,970 problem instances with varying demand profiles.
- Architecture: Includes three convolutional layers followed by max-pooling layers, and two fully connected layers for multi-label classification.
  - First layer: 32 filters, kernel size 10.
  - Second layer: 64 filters, kernel size 5.
  - Third layer: 128 filters, kernel size 3.
  - Fully connected layers: 256 neurons and 128 neurons with dropout regularisation.
  - Output layer
- Training uses the Adam optimiser with binary cross-entropy loss for classification accuracy

Current Research & Techniques

AI/ML Integrated Optimisation

The reduced MILP model achieves:

- Up to 81% reduction in constraints and 83% reduction in binary variables compared to the original model.
- Infeasibility or sub-optimality in 10.65% of scenarios, highlighting areas for improvement in prediction accuracy.

Conclusions & Future Prospects

- We explored some ways to tackle complexity of MILP models
- Modern or future computer architectures can be an opportunity to push further the advances in MILP-based algorithms (i.e., GPU architecture and distributed systems)
- Exploiting generative AI to produce MILP models from a textual description of the problem or examples of data and solutions. An. early research has been done in (AhmadiTeshnizi, Gao, & Udell, 2023). Although it may not impact the work of MILP experts, this may change how end-users can access optimisation tools.
- Pushing solvers to the next level

References

- Müllerklein, D., & Fontaine, P. (2025). Resilient transportation network design with disruption uncertainty and lead times. European Journal of Operational Research, 322(3), 827-840.
- Clautiaux, F., & Ljubić, I. (2024). Last fifty years of integer linear programming: a focus on recent practical advances. European Journal of Operational Research.
- Bigi, G., Frangioni, A., Gallo, G., Pallottino, S., & Scutellà, M. G. (2024). Appunti di Ricerca Operativa. Dipartimento di Informatica, Università di Pisa. Retrieved March 20, 2025, from https://commalab.di.unipi.it/wp-content/uploads/2024/05/v1.0.0-240424.pdf.
- Triantafyllou, N., & Papathanasiou, M. M. (2024). Improved Deep Learning Architectures for the Decomposition of Mixed Integer Optimization Problems. In Computer Aided Chemical Engineering (Vol. 53, pp. 2749-2754). Elsevier.
- Lanza, G., Passacantando, M., & Scutellà, M. G. (2023). Matheuristic approaches to the green sequencing and routing problem. Flexible Services and Manufacturing Journal, 1-52.
- AhmadiTeshnizi, A., Gao, W., & Udell, M. (2023). Optimus: Optimization modeling using mip solvers and large language models. arXiv preprint arXiv:2310.06116.
- Hillier, F. S., & Lieberman, G. J. (2015). Chapter 2. Introduction to operations research. McGraw-Hill.
- Benders, J. (1962). Partitioning procedures for solving mixed-variables programming problems. Numer. Math, 4(1), 238-252.

Thank you for your attention!

Let us Discuss!

How do you think we can master MILP?
Can deep learning replace traditional heuristics?